AFRL-SN-WP-TP-2003-105

NAVIGATION PATTERNS AND USABILITY OF ZOOMABLE USER INTERFACES WITH AND WITHOUT AN OVERVIEW

Kasper Hornbæk
Benjamin B. Bederson
Catherine Plaisant

DECEMBER 2002

Approved for public release; distribution is unlimited.

SENSORS DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7318

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

1. REPORT DATE ***(DD-MM-YY)***
December 2002

2. REPORT TYPE
Journal Article

3. DATES COVERED ***(From - To)***

4. TITLE AND SUBTITLE
NAVIGATION PATTERNS AND USABILITY OF ZOOMABLE USER INTERFACES WITH AND WITHOUT AN OVERVIEW

5a. CONTRACT NUMBER
F33615-97-1-1018

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
62301E

5d. PROJECT NUMBER
ARPA

5e. TASK NUMBER
AA

5f. WORK UNIT NUMBER
1P

6. AUTHOR(S)
Kasper Hornbæk (University of Copenhagen)
Benjamin B. Bederson and Catherine Plaisant (University of Maryland)

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Maryland
Department of Computer Science
Human-Computer Interaction Laboratory
College Park, MD 20742

University of Copenhagen

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Sensors Directorate
Air Force Research Laboratory
Air Force Materiel Command
Wright-Patterson Air Force Base, OH 45433-7318

10. SPONSORING/MONITORING AGENCY ACRONYM(S)
AFRL/SNAR

11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)
AFRL-SN-WP-TP-2003-105

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
Published in *ACM Transactions on Computer-Human Interaction*, Vol. 9, No. 4, December 2002, pp. 362-389.

See other published work in the DTIC collection under contract number F33615-97-1-1018.

14. ABSTRACT ***(Maximum 200 Words)***
The literature on information visualization establishes the usability of interfaces with an overview of the information space, but for zoomable user interfaces, results are mixed. We compare zoomable user interfaces with and without an overview to understand the navigation patterns and usability of these interfaces. Thirty-two subjects solved navigation and browsing tasks on two maps. We found no difference between interfaces in subjects' ability to solve tasks correctly. Eighty percent of the subjects preferred the interface with an overview, stating that it supported navigation and helped keep track of their position on the map. However, subjects were faster with the interface without an overview when using one of the two maps. We conjecture that this difference was due to the organization of that map in multiple levels, which rendered the overview unnecessary by providing richer navigation cues through semantic zooming. The combination of that map and the interface without an overview also improved subjects' recall of objects on the map. Subjects who switched between the overview and the detail windows used more time, suggesting that integration of overview and detail windows adds complexity and requires additional mental and motor effort.

15. SUBJECT TERMS
Experimentation, Human Factors, Measurement, Performance, Information visualization, zoomable user interfaces (ZUIs), overviews, overview+detail interfaces, navigation, usability, maps, levels of detail

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT:**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** (Monitor)
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**	SAR	34	Jason Johnson
Unclassified	Unclassified	Unclassified			**19b. TELEPHONE NUMBER** *(Include Area Code)* (937) 255-5668 x4047

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39-18

Navigation Patterns and Usability of Zoomable User Interfaces with and without an Overview

KASPER HORNBÆK
University of Copenhagen
and
BENJAMIN B. BEDERSON and CATHERINE PLAISANT
University of Maryland

The literature on information visualization establishes the usability of interfaces with an overview of the information space, but for zoomable user interfaces, results are mixed. We compare zoomable user interfaces with and without an overview to understand the navigation patterns and usability of these interfaces. Thirty-two subjects solved navigation and browsing tasks on two maps. We found no difference between interfaces in subjects' ability to solve tasks correctly. Eighty percent of the subjects preferred the interface with an overview, stating that it supported navigation and helped keep track of their position on the map. However, subjects were faster with the interface without an overview when using one of the two maps. We conjecture that this difference was due to the organization of that map in multiple levels, which rendered the overview unnecessary by providing richer navigation cues through semantic zooming. The combination of that map and the interface without an overview also improved subjects' recall of objects on the map. Subjects who switched between the overview and the detail windows used more time, suggesting that integration of overview and detail windows adds complexity and requires additional mental and motor effort.

Categories and Subject Descriptors: H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*evaluation/methodology; interaction styles (e.g., commands, menus, forms, direct manipulation)*; I.3.6 [**Computer Graphics**]: Methodology and Techniques—*interaction techniques*

General Terms: Experimentation, Human Factors, Measurement, Performance

Additional Key Words and Phrases: Information visualization, zoomable user interfaces (ZUIs), overviews, overview+detail interfaces, navigation, usability, maps, levels of detail

1. INTRODUCTION

Information visualization [Card et al. 1999] has become a successful paradigm for human-computer interaction. Numerous interface techniques have been

This work was funded in part by DARPA's Command Post of the Future project, contract number F336159711018, and ChevronTexaco.
Authors' addresses: K. Hornbæk, Department of Computing, University of Copenhagen, Universitetsparken 1, DK-2100 Copenhagen Ø, Denmark; email: kash@diku.dk; B. B. Bederson and C. Plaisant, Department of Computer Science, Human-Computer Interaction Laboratory, University of Maryland, College Park, MD 20742; email: {bederson,plaisant}@cs.umd.edu.

ACM Transactions on Computer-Human Interaction, Vol. 9, No. 4, December 2002, Pages 362–389.

proposed and an increasing number of empirical studies describe the benefits and problems of information visualization, for example, Beard and Walker [1990], Schaffer et al. [1996], Hornbæk and Frøkjær [1999], Chen and Czerwinski [2000]. Interfaces with an overview and zoomable user interfaces have been extensively discussed in the literature on information visualization. Interfaces with an overview, often called *overview+detail interfaces* [Plaisant et al. 1995], show the details of an information space together with an overview of the entire information space. Such interfaces can improve subjective satisfaction (e.g., North and Shneiderman [2000]), and efficiency (e.g., Beard and Walker [1990]). Zoomable user interfaces organize information in space and scale, and use panning and zooming as their main interaction techniques [Perlin and Fox 1993; Bederson et al. 1996]. Research prototypes of zoomable user interfaces include interfaces for storytelling [Druin et al. 1997], Web browsing [Hightower et al. 1998], and browsing of images [Combs and Bederson 1999; Bederson 2001]. However, few empirical studies have investigated the usability of zoomable user interfaces, and the results of those studies have been inconclusive. In addition, the usability of overviews for zoomable user interfaces has not been studied.

In this article we present an empirical analysis of zoomable user interfaces with and without an overview. We investigate the following:

—how the presence or absence of an overview affects usability;

—how an overview influences the way users navigate information spaces; and

—how different organizations of information spaces may influence navigation patterns and usability.

With this work, we aim to strengthen the empirical literature on zoomable user interfaces, thereby identifying challenges for researchers and advising designers of user interfaces.

In Section 2, we review the literature on overviews and zoomable user interfaces. Then, we present our empirical investigation of differences in navigation patterns and usability in zoomable user interfaces with and without an overview. Finally, we discuss the trade-off between time and satisfaction in such interfaces and explain the interaction between usability and differently organized information spaces.

2. RELATED WORK

This section summarizes the research questions and empirical findings about interfaces with overviews and zoomable user interfaces. It explains the literature behind our design decisions and the motivation for the experiment, both described in subsequent sections.

2.1 Interfaces with Overviews

Interfaces with overviews present multiple views of an information space where some views show detailed information about the information space (called *detail windows*), while other views show an overview of the information space (called *overview windows* or *overviews*). Examples of such interfaces include editors

for program code [Eick et al. 1992], interfaces for image collections [North et al. 1995], and commercial programs such as Adobe Photoshop.[1] Interfaces with an overview have been found to have three benefits. First, navigation is more efficient because users may navigate using the overview window rather than using the detail window [Beard and Walker 1990]. Second, the overview window aids users in keeping track of their current position in the information space [Plaisant et al. 1995]. The overview window itself might also give users task-relevant information, for example, by enabling users to read section titles from an overview of a document [Hornbæk and Frøkjær 2001]. Third, the overview gives users a feeling of control [Shneiderman 1998]. A drawback of interfaces with an overview is that the spatially indirect relation between overview and detail windows might strain memory and increase the time used for visual search [Card et al. 1999, p. 307]. In addition, such interfaces require more screen space than interfaces without overviews.

Taxonomies and design guidelines for overviews [Beard and Walker 1990; Plaisant et al. 1995; Carr et al. 1998; Baldonado et al. 2000] contain three main points. First, the overview and detail windows need to be tightly coupled [Ahlberg and Shneiderman 1994], so that navigation or selection of information objects in one window is immediately reflected in the other windows. Tight coupling of overview and detail views has been found useful in several studies (e.g., North and Shneiderman [2000]). Second, for any relation between overview and detail windows, the zoom factor is the ratio between the larger and smaller of the magnification of the two windows. For overview+detail interfaces, this factor is recommended to be below 25 [Plaisant et al. 1995] or below 30 [Shneiderman 1998]. It is unclear, however, if the sizes of the detail and overview windows influence the recommended zoom factor. Third, the size of the overview window influences how much information can be seen at the overview and how easy it is to navigate on the overview. However, a large overview window might take screen real estate from the detail window. Plaisant et al. [1995] argued that the most usable sizes of the overview and detail windows are task dependent. A large overview window, for example, is required for a monitoring task, while a diagnostic task might benefit from a large detail window.

A number of empirical studies have found that having an overview improves user satisfaction and efficiency over interfaces without an overview. Beard and Walker [1990] compared the effect of having an overview window to navigating with scrollbars. In a 280-word ordered tree, subjects used an overview window that allowed dragging a field-of-view and one that allowed both dragging and resizing the field-of-view. For tasks where subjects tried to locate a word in the tree and tasks where they repeatedly went from one side of the tree to the other, the overview window led to significantly faster task completion. North and Shneiderman [2000] compared 18 subjects' performance with a detail-only, an uncoordinated overview+detail, and a coordinated overview+detail interface for browsing textual population data. Compared to the detail-only interface, the coordinated interface was 30–80% faster and scored significantly higher on a satisfaction questionnaire. Hornbæk and Frøkjær [2001] compared an interface

[1]See http://www.adobe.com/photoshop/.

with an overview for electronic documents to a fisheye and a detail-only interface. Essays produced with aid of the interface with an overview scored significantly higher than essays produced with the detail-only interface. However, for tasks that required subjects to answer a specific question, the interface with an overview was 20% slower compared to the detail-only interface. All but one of the 21 subjects preferred having the overview.

2.2 Zoomable User Interfaces

While zoomable user interfaces have been discussed since at least 1993 [Perlin and Fox 1993], no definition of zoomable user interface has been generally agreed upon. In this article, we consider the two main characteristics of zoomable user interfaces to be (a) that information objects are organized in space and scale, and (b) that users interact directly with the information space, mainly through panning and zooming. In zoomable user interfaces, space and scale are the fundamental means of organizing information [Perlin and Fox 1993; Furnas and Bederson 1995]. The appearances of information objects are based on the scale at which they are shown. Most common is geometric zoom, where the scale linearly determines the apparent size of the object. Objects may also have a more complex relation between appearance and scale, as in so-called semantic zooming [Perlin and Fox 1993; Frank and Timpf 1994], which is supported in the zoomable user interface toolkit Jazz [Bederson et al. 2000]. Semantic zooming is commonly used with maps, where the same area on the map might be shown with different features and amounts of detail depending on the scale. Constant density zooming [Woodruff et al. 1998a] introduces a more complex relation between scale and appearance where the number of objects currently shown controls the appearance of objects, so that only a constant number of objects is visible simultaneously.

The second main characteristic of zoomable user interfaces is that the information space is directly visible and manipulable through panning and zooming. Panning changes the area of the information space that is visible, and zooming changes the scale at which the information space is viewed. Usually, panning and zooming are controlled with the mouse or the keyboard, so that a change in the input device is linearly related to how much is panned or zoomed. Nonlinear panning and zooming have been proposed in three forms: (a) goal-directed zoom, where direct zooming to an appropriate scale is supported [Woodruff et al. 1998b]; (b) combined zooming and panning, where extensive panning automatically leads to zooming [Igarishi and Hinckley 2000]; and (c) automatic zoom to objects, where a click with the mouse on a object automatically zooms to center on that object [Furnas and Zhang 1998; Ware 2000]. When zooming, two ways of changing scale are commonly used. In jump zooming, the change in scale occurs instantly, without a smooth transition. Jump zooming is used in Pad [Perlin and Fox 1993], Schaffer et al.'s [1996] experimental system, and commercial systems such as Adobe PhotoShop or MapQuest.[2] In animated zooming the transition from the old to the new scale is smooth [Bederson and Hollan 1994; Pook et al. 2000; Bederson et al. 2000]. An important issue in animated zooming is the

[2]See http://www.mapquest.com/.

duration of the transition and the user's control over the zooming speed, that is, the ratio between the zooming time and the zooming factor. Guo et al. [2000] provided preliminary evidence that a zoom speed around 8 factors/s is optimal. Card et al. [1991] argued that the zoom time should be approximately 1 s, although in some zoomable user interfaces, for example, Jazz, users can control both the zoom time and the zoom factor. Bederson and Boltman [1999] investigated whether an animated or jump zoom technique affected 20 subjects' ability to remember the topology of and answer questions about a nine-item family tree. Subjects were better at reconstructing the topology of the tree using animated zooming, but no difference in satisfaction or task completion time was found.

The empirical investigations of zoomable user interfaces are few and inconclusive. Páez et al. [1996] compared a zoomable user interface based on Pad++ [Bederson and Hollan 1994] to a hypertext interface. Both interfaces gave access to a 9-page scientific paper. In the zoomable user interface, the scale of the sections and subsections of the paper were manipulated, so that the entire paper fit on the initial screen. No significant difference was found between the two interfaces for the 36 subjects' satisfaction, memory for the text, or task completion time. Schaffer et al. [1996] compared 20 subjects' performance with a zoomable user interface and a fisheye interface. Subjects had to locate a broken link in a telephone network and reroute the network around the link. Subjects used 58% more time for completing the task in the zoomable user interface. Subjects seemed to prefer the fisheye interface, although this was not clearly described in the paper.

Hightower et al. [1998] presented two experiments that compared the history mechanism in Netscape Navigator with a graphical history in a zoomable user interface called *PadPrints*. In the first experiment, 37 subjects were required to answer questions about Web pages. No significant difference in task completion time was found, but subjects preferred the PadPrints interface. In the second experiment, subjects were required to return to already visited Web pages. Subjects were approximately 40% faster using the PadPrints interface and preferred PadPrints to Netscape Navigator. Combs and Bederson [1999] compared four image browsers: two commercial 3D interfaces, one commercial 2D interface, and an image browser based on Pad++. Thirty subjects searched for images in an image database that they had just browsed. Subjects were significantly faster using the 2D and the zoomable user interfaces, especially as the number of images in the database went from 25 to 225. The study presented some evidence that recall of images is improved in the zoomable user interface, but found no difference in subjective satisfaction between interfaces. Ghosh and Shneiderman [1999] compared 14 subjects' use of an overview+detail and a zoomable user interface to personal histories, LifeLines [Plaisant et al. 1996]. The zoomable user interface was marginally slower than the overview+detail interface. No difference in subjective satisfaction was found.

In general, the experimental results about zoomable user interfaces are mixed, reflecting differences in the interfaces that zoomable user interfaces are compared to, in the organization and size of the information spaces used, and in the implementation of zooming. In addition, the characteristics of zoomable user

interfaces and interfaces with an overview are increasingly blended. For example, zoomable user interfaces have been combined with transparent overviews [Pook et al. 2000]; some interfaces with overviews have been extended with animated zooming [Ghosh and Shneiderman 1999]; and some effort has been put into extending zoomable user interfaces with navigation mechanisms that supplement zooming and panning (see, for example, Jul and Furnas [1998]). The main difference between research in zoomable user interfaces and in interfaces with an overview is that research in zoomable user interfaces has investigated the usefulness of zooming as a way of navigating, while other research has focused on the impact of a coupled overview. As interfaces with an overview begin to use panning and zooming as their main navigation technique and as zoomable user interfaces begin to provide overviews and other navigation aids, the central research questions become (1) what is the difference between different techniques for controlling and executing zooming, possibly taking into account the presence of an overview and other navigation aids; and (2) what is the effect of an overview (or other navigation aids), given that the interface provides pan and zoom techniques. In the experiment presented next, we address the latter question.

3. EXPERIMENT

To understand the differences in navigation patterns and usability between zoomable user interfaces with and without an overview, we conducted a controlled experiment. In the experiment, subjects used interfaces we will call the *overview interface* and *no-overview interface* to solve 10 tasks on each of two differently organized maps.

3.1 Hypotheses

In addition to the three aims mentioned in the introduction, three hypotheses guided the design of the experiment:

(1) Recall of objects on the map would be better in the no-overview interface. Zoomable user interfaces have been speculated to improve understanding of large information spaces, because of the integrated experience of the information space [Furnas and Bederson 1995]. As mentioned in Section 2, one experiment [Combs and Bederson 1999] found improved recall in zoomable user interfaces. In the interface with an overview, we expected subjects to occasionally use the overview window for navigation in the overview+detail interface, thereby losing the integrated experience of the information space. In addition, research has shown that users have difficulty in integrating multiple views [Card et al. 1999, p. 634]; lower recall with the overview interface may be one measurable implication of these observations.

(2) Subjects would prefer the overview interface, because of the information contained on the overview window and the additional navigation features. This hypothesis was based on the research on nonzoomable interfaces with overviews, summarized in Section 2.

(3) The overview interface would be faster for tasks that require comparison of information objects and scanning large areas (the latter we called *browsing tasks*). The literature suggests that comparison and scanning tasks are particularly well supported by an overview because the overview can be used for jumping between objects to be compared and because it can help subjects to keep track of which parts of the information space have already been explored.

3.2 Subjects

Thirty-two subjects participated in the experiment, 23 males and 9 females. Subjects were recruited at the University of Maryland and received 15 US dollars for participating in the experiment. The age of the subjects ranged from 18 to 38; the mean age was 23.4 years. Twenty-three subjects were computer science or engineering students, four had other majors, and five were research staff or loosely affiliated with the university. Thirty-one subjects used computers every day. Twenty-three of the subjects had never used zoomable user interfaces, while nine subjects had seen or used a zoomable user interface prior to participating in the experiment. We required that subjects had spent less than 2 weeks in the states of Washington and Montana, because the experiment used maps of those states.

3.3 Interfaces

For the experiment, we constructed an overview and a no-overview interface, both based on the zoomable user interface toolkit Jazz [Bederson et al. 2000]. When users held down the left mouse button, zooming in began after a delay of 400 ms. Users zoomed out by holding down the right mouse button. The maximum zoom factor was 20, meaning that subjects could view the map at scale 1 through scale 20. At scale 1, the initial unmagnified view of the map was shown; at scale 20 the initial view of the map was magnified 20 times. The zoom speed was 8 factors/s; that is, subjects could zoom from the initial view of the map to the maximum magnification in 2.5 s. Users panned by holding down the left mouse button and moving the mouse in the opposite direction of what they wished to see (i.e., the map followed the mouse). In the lower right corner of both interfaces was an icon showing the four compass points, which were referred to in some tasks. Next to this icon was a button labeled *zoom out*, which when pressed zoomed out to the initial view of the map. This button was expected to help subjects return to the initial view of the map if they were lost.

The no-overview interface is shown in Figure 1. Subjects could only interact with this interface using the zoom and pan techniques described above.

The overview interface is shown in Figure 2. In the top-right corner of the interface, an overview window shows the entire map at one-sixteenth the size of the detail window. This choice was arbitrary, lacking design guidelines on overview sizes (see Section 2.1). However, it was similar to the average size of the overviews we were familiar with. The current location of the detail window on the map was indicated in the overview window by a 70% transparent field-of-view box. The overview and detail windows were tightly coupled, so that

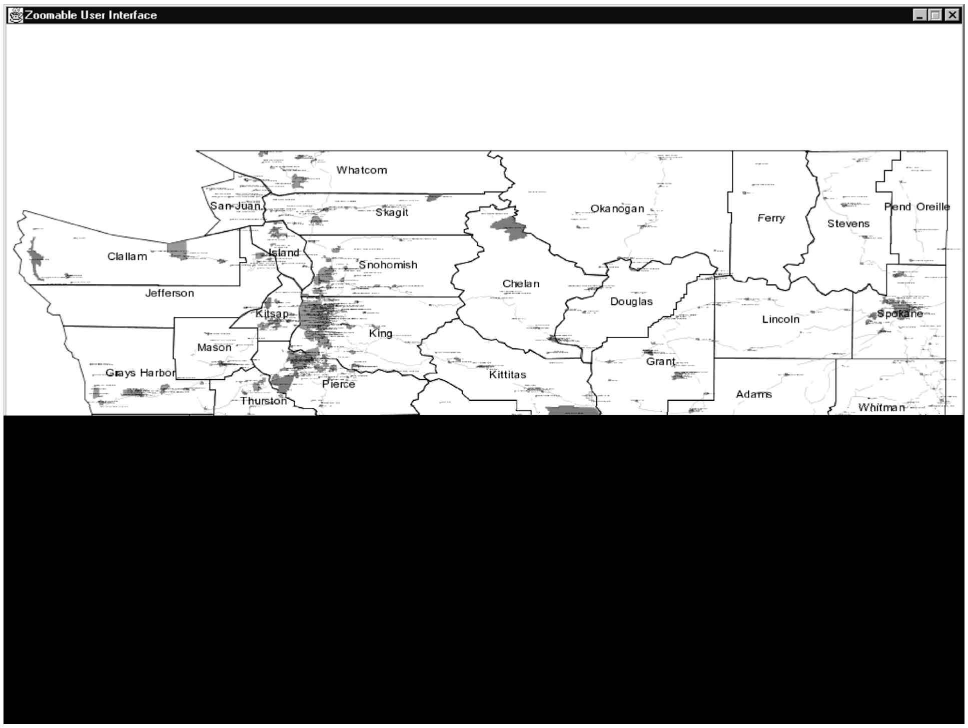

Fig. 1. *No-overview interface showing the Washington map*. The user may zoom and pan to change the area of the map shown. In the lower right corner of the window a button is shown that will zoom out to the initial view of the map. Next to this button is an indication of the four compass points. The colors of the map are reproduced here as different shades of gray. The map is shown at scale 1, that is, at the initial view of the map.

zooming or panning in the detail window immediately updated the overview window and dragging the field-of-view box changed which part of the map was shown in the detail window. The subjects could also click in the overview window outside of the field-of-view box, which centered the field-of-view box on the point clicked on. The field-of-view box could be resized by dragging the resize handle in the bottom right corner of the field-of-view box. The subjects could also draw a new field-of-view box by holding down the left button and moving the mouse until the desired rectangle was drawn. The field-of-view box always kept the same aspect ratio, which corresponded to the detail window and the overview window.

3.4 Maps

The motivation for using maps for the experiment was threefold. First, interfaces for maps constitute an important area of research. Second, maps include characteristics of other, commonly used information structures, for example, hierarchical information (nesting of information objects) and network information (connections between information objects). Therefore, results concerning maps may be generalized to other information structures. Third, the direct relation

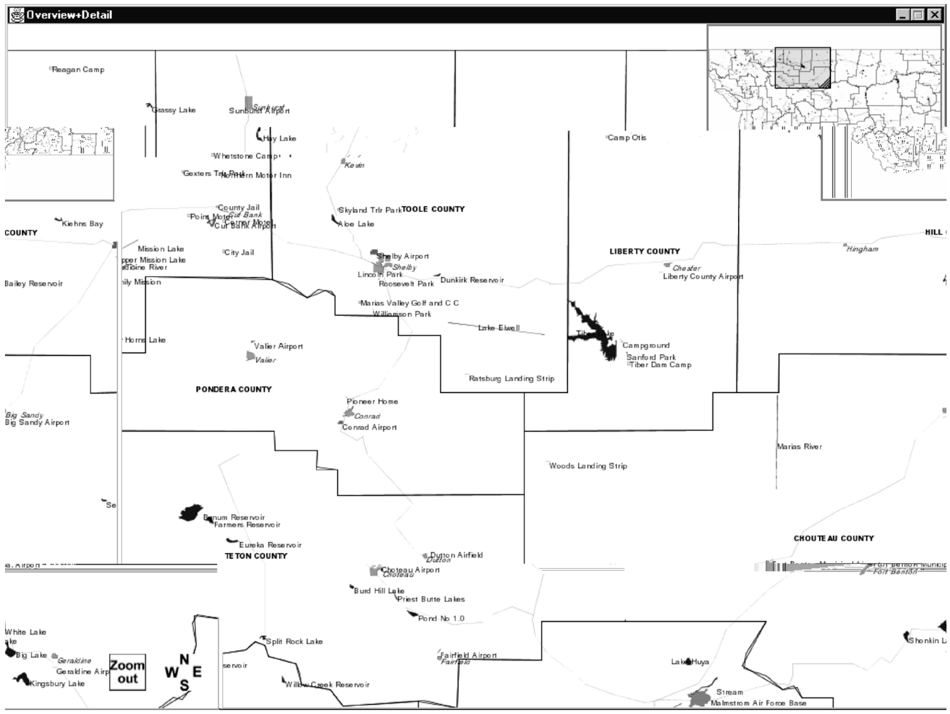

Fig. 2. *The overview interface showing the Montana map*. In the top right corner of the interface is the overview window, which shows an overview of the entire map. The gray area in the overview window is the field-of-view box that indicates which part of the map is currently shown in the detail window. In the bottom right corner of the field-of-view box is the resize handle that allows the user to make the field-of-view smaller or larger, that is, to zoom in or out. The two buttons in the lower right corner are similar to the buttons in the zoomable user interface. The map is shown at scale 4, meaning that the objects in the detail window are magnified 4 times.

between representation and physical reality aids interpretation of maps compared to the often difficult interpretation of abstract information spaces [Hornbæk and Frøkjær 1999].

We created two maps based on data from the 1995 United States Census.[3] The maps contained eight types of map objects: counties, cities, parks, airports, lakes, railroads, military installations, and other landmarks. Each map object, except railroads, consisted of a shape and a label. A distinct color identified each type of map object. In addition, county names were shown in bold type and city names in italic type. The maps were organized by placing labels for map objects at different scales, changing the apparent size of the labels as follows (also see Figure 3):

—The map of the state of Washington showed map objects at three levels of scale: county level (scale 1, 39 labels), city level (scale 5, 261 labels), and landmark level (scale 10, 533 labels). At the county level, labels were the same size as a 10-point font when the map was zoomed out (i.e., at scale 1)

[3]See http://www.census.gov/geo/www/tiger/ or http://www.esri.com/data/online/tiger/.

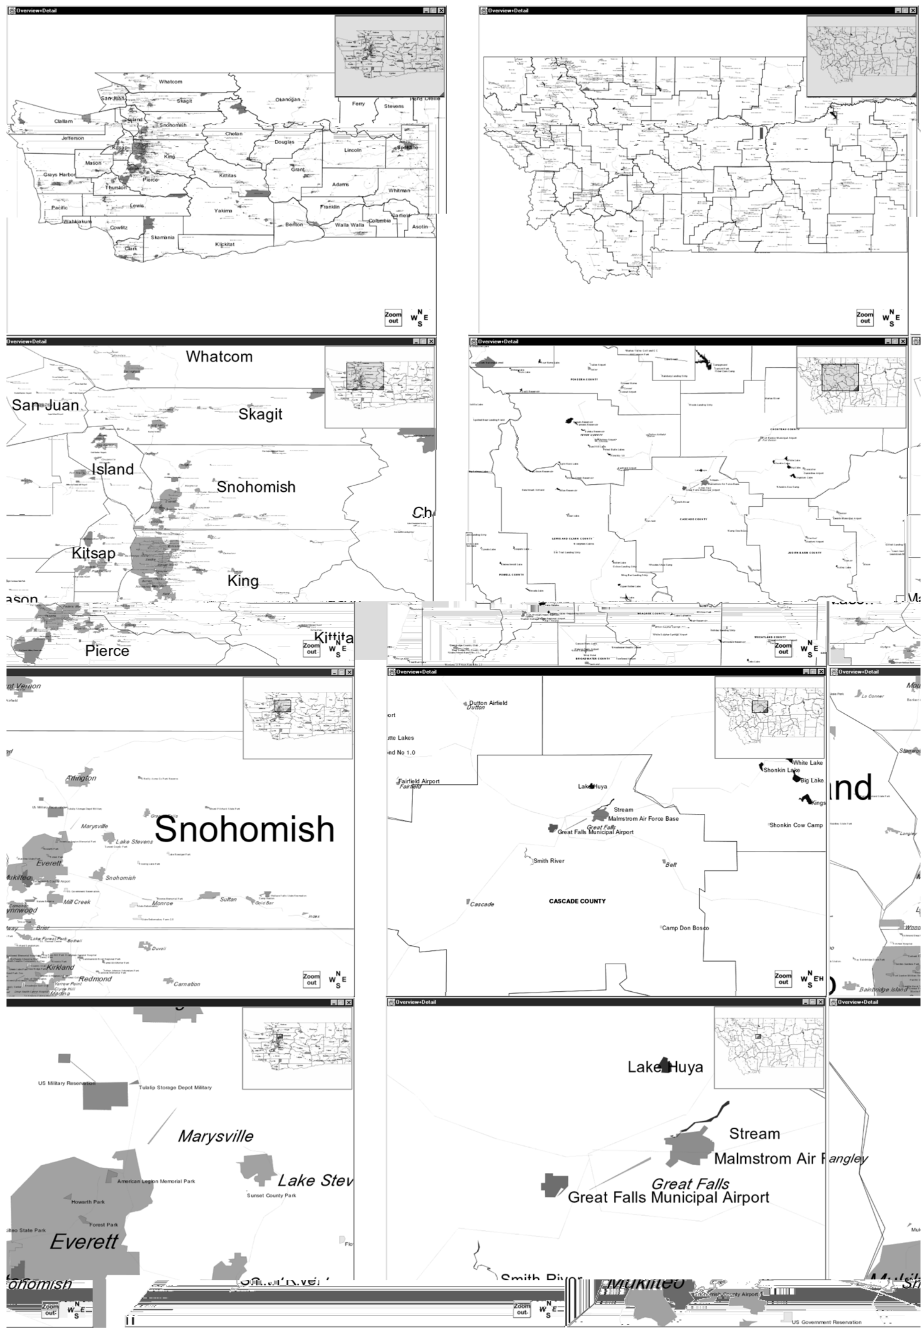

Montana map **Washington map**

Fig. 3. *Eight screenshots of the maps*. The four screenshots in the left column show the Washington map; the right column shows the Montana map. From top to bottom the maps are shown at scales 1, 3, 7, and 20. On the Washington map, map objects are labeled at three different levels: county level (39 counties, for example, Snohomish in the left column, screenshot 2 from the top), city level

and larger when the map was magnified. When labels were shown at the

We also gave the subjects two recall tasks that test their memory of the structure and content of the map. The first recall task consisted of five small maps showing the outline of the state depicted on the map. For three of these small maps, a part of the map was darkened and the subjects were asked to write down as many objects within the dark area as they remembered. For two of the maps, subjects themselves could mark a county on the map with a cross, and write down any map objects they remembered within that county. The second recall task consisted of three county names, each associated with a list of 10 cities. Subjects were told to circle all cities within a county and cross out cities they were confident were not located in the county mentioned. The list of cities consisted of the three largest cities within the county mentioned, the three largest cities in counties just next to the county mentioned, and four cities in entirely different areas of the map.

3.6 Experimental Design and Dependent Variables

The experiment varied interface type (no-overview vs. overview), task type (navigation vs. browsing tasks), and map (Washington vs. Montana map), in a within-subjects balanced factorial design. The experiment consisted of two parts. In the first part, subjects used one interface giving access to one map and performed five navigation and five browsing tasks. In the second part, subjects used the other interface in combination with the not-yet explored map. Subjects were randomly assigned to one of the four possible combinations of interface and map. Within each of these four combinations, subjects were further randomly assigned to one of four permutations of task types in the two parts. Each of the resulting 16 groups contained two subjects. The order of the five tasks within a task type was the same for all subjects.

We used a range of dependent variables to capture information about navigation patterns and usability:

—*Accuracy in answering questions.* Accuracy was calculated as the number of answers that were correct (all map objects given as answer to a task correct), partially correct (one correct and one wrong map object), and wrong (all map objects wrong).
—*Recall of map objects.* For the recall task that required subjects to mark counties and cities on the map, we counted as correct the number of counties and cities within 1 cm of the actual location. For the recall task that required subjects to recognize the cities in a county after they had finished using the interface, we measured the number of correct indications, corrected with a penalty for guessing (the number of wrong guesses divided by the number of wrong answer possibilities for the question).
—*Task completion time.* Task completion time was measured as the time subjects could see the map. The time subjects used for the initial reading of the task, as well as the time used for entering answers, was not included.
—*Preference.* Preference was determined from subjects' indication of which interface they preferred using and from the reasons subjects gave for their indication.

—*Satisfaction.* Satisfaction was measured using seven questions with 9-point semantic differentials. Five of the questions were taken from the Questionnaire for User Satisfaction [Chin et al. 1988] and two questions were custom made. The wording of the questions appears in Figure 6.

—*Navigation actions.* We logged all interactions with the interfaces and measured the number of pan actions in the detail window and on the overview window (centering or dragging the field-of-view). We also measured zoom actions in the detail window and on the overview (resizing the field-of-view). An action was initiated when the mouse button signifying that action was pressed and was ended either when the button was released or when more than 1 s passed without any logged mouse movements. To compare these measures across interfaces, we combined them into a measure of total distance panned and the sum of scale changes, that is, the amount zoomed.

3.7 Procedure

The interfaces were run on a 650-MHz Pentium III laptop with an ordinary mouse. The screen was 13 in, with a resolution of 1024×768.

Upon arriving at the lab, subjects filled out a questionnaire about gender, occupation, and familiarity with computers. Then, subjects were introduced to the two interfaces and tried three practice tasks that lasted on average 11 min.

The main phase of the experiment consisted of two parts, each containing 10 tasks. For each task, subjects initially saw a window that covered the entire map. After reading a piece of paper that described the task, subjects clicked on a button to see a zoomed-out view of the map. When subjects had completed the task, they entered their answer using a tightly coupled text field and list box containing the labels of all objects on the map. For all tasks, subjects were asked to proceed to the next task when they had searched for 5 min. After solving all tasks in the first part of the experiment, subjects received the recall tasks and filled out a satisfaction questionnaire about the interface just used. After a 5-min break, subjects began the second part of the experiment, which used the same procedure as the first part.

After the second part of the experiment, subjects filled out a form about which interface they preferred. On average, the experiment lasted 1.5 h.

4. RESULTS

In Sections 4.1 through 4.4, we use univariate analyses of variance to investigate the accuracy of answers to tasks, recall of map objects, preference and satisfaction, and how subjects navigated. In the analyses, the four possible combinations of interface and map type are a between-subject factor (3° of freedoms, or df). This leaves 28 df for the error term. Within-subject factors are interface (1 df), map type (1 df), task type (1 df), and tasks nested within map type and task type (4 df). For each dependent variable, these factors and their interactions are used as the model for the analysis of variance. For the dependent variables satisfaction and recall, only interface and map type are used, as these variables were only measured once in each part of the experiment.

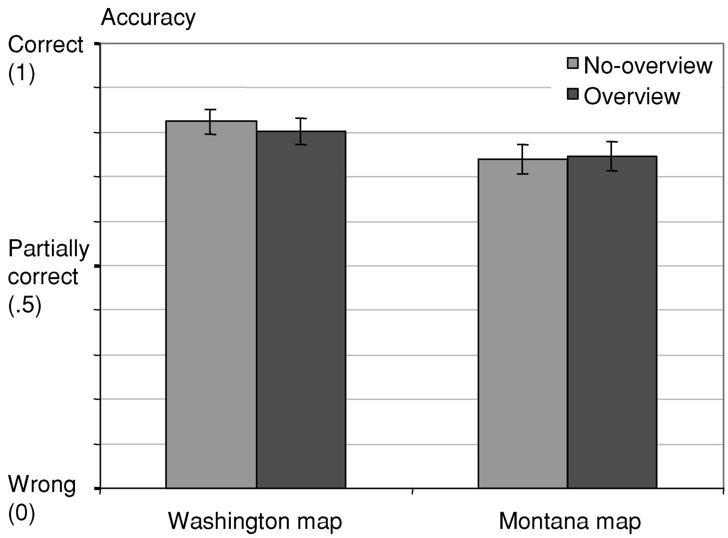

Fig. 4. *The average accuracy for the answers to the experimental tasks (for each bar $N = 160$).* The figure shows the average accuracy for the two interfaces between maps. The answers to each task were scored as 1 for correct, .5 for partially correct, and 0 for wrong. A partially correct answer mentioned only one out of two map objects correctly. Error bars show the standard error of the mean [SD/SQRT(N)].

4.1 Accuracy and Recall

Figure 4 summarizes the accuracy of the answers to the experimental tasks. We found no difference in the accuracy between interfaces, $F(1, 28) = .144$, $p > .5$. Between the two maps, a significant difference in the number of tasks correctly answered can be found, $F(1, 28) = 11.63$, $p < .001$. Tasks solved on the Washington map were more often answered correctly than tasks solved on the Montana map.

Figure 5 shows the measures of recall of map objects for the two interfaces. With the overview interface, subjects did better at the recall task with the Montana map compared to the Washington map. The no-overview interface showed the opposite pattern. These patterns were confirmed with a rank-based test of the number of marked cities and counties by a significant interaction between interface and map type, $F(1, 28) = 4.25$, $p < .05$. No such interaction was found for the number of recognized cities, $F(1, 28) = 1.69$, $p > .2$; only a marginally significant difference between interfaces for the Washington map was found, $F(1, 28) = 3.90$, $p < .06$.

Large individual differences existed in the accuracy and recall of cities and counties. One subject correctly answered 19 of the 20 questions; another subject answered only nine questions correctly. In the recall task, one subject marked on average 11 cities or counties on the map; another subject marked none.

4.2 Preference and Satisfaction

Twenty-six subjects stated that they preferred using the overview interface, while six subjects stated they preferred the no-overview interface. Thus,

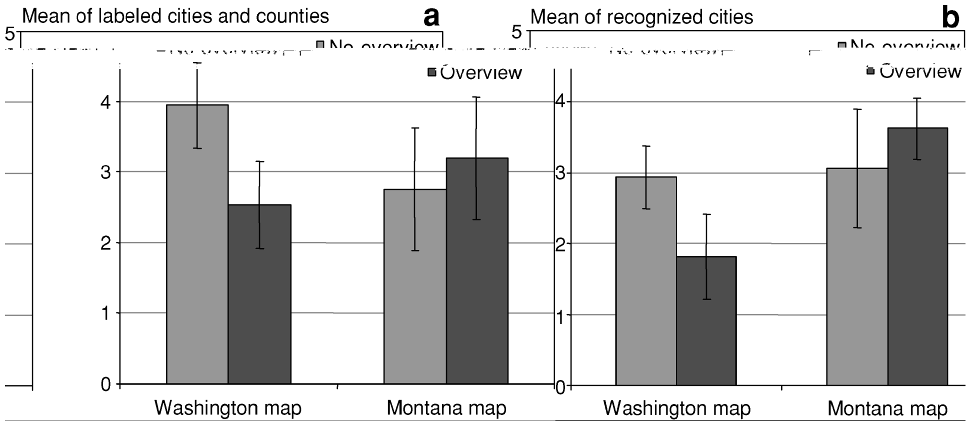

Fig. 5. *Mean number of correct answers to recall tasks* (*for each bar* $N = 16$). Panel a shows the mean number of correctly labeled cities and counties for each subject in one of the two parts of the experiment; panel b shows the mean number of correctly recognized cities, adjusted for guessing. Error bars show the standard error of the mean.

significantly more subjects preferred the overview interface, $\chi^2(1, N = 32) = 12.5$, $p < .001$. Subjects explained their preference for the overview interface as follows:

—The overview window provided information about the current position on the map; for example, one subject wrote: "It is easier to keep track of where I am." $N = 9$ subjects made similar comments.
—The overview window supported navigation ($N = 7$); for example, one subject wrote: "[It was] easier to navigate in the overview box while looking at the detail map for answers." Two subjects wrote similar comments at the end of the part of the experiment in which they had used the overview+detail interface.
—The overview window was helpful when scanning a large area ($N = 4$); for example, one subject wrote: "It made surveying a large map less disorienting especially when small landmarks had to be spotted."
—The overview window was useful for zooming ($N = 2$); for example, one subject wrote: "The zoom feature in the top right was extremely helpful."
—The overview window supported comparing objects ($N = 2$); for example, one subject worte: "Easier to move between counties while at the same zoom level -> easier to compare the size of objects".

The six subjects who preferred the no-overview interface mentioned the following:

—Locating objects felt faster using the no-overview interface ($N = 2$); for example, one subject wrote: "I found myself answering my tasks much quicker using the [no overview] interface."
—One subject preferred the no-overview interface because the overview window got in the way when using the overview interface: "Overview+detail

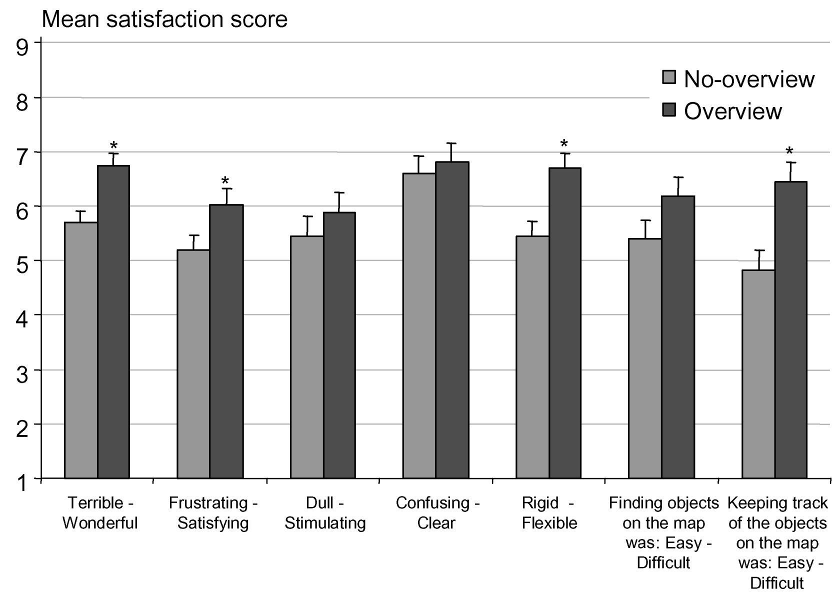

Fig. 6. *Satisfaction with the interfaces* (*for each bar* $N = 32$). The figure shows the mean score for the seven satisfaction questions that each subject answered after using each of the two interfaces. Error bars indicate the standard error of the mean. The questions were answered on a 9-point semantic differential going from 1 (lowest score) to 9 (highest score). Significant differences at the .05 level are marked in the figure with an asterisk (*).

would seem to be more powerful, but the abundance of features got in the way to the effect of imposing on usability." Three subjects made similar comments at the end of the part of the experiment where they used the overview+detail interface. Nevertheless, these subjects preferred the overview interface.

In addition, four subjects commented that they found it hard to resize the field-of-view box; three subjects commented that the map seemed larger using the no-overview interface; two subjects commented that when using the no-overview interface it was sometimes unclear where they were on the map; and two subjects commented that it was useful that the overview window gave a visual indication of the current zoom factor.

Figure 6 shows the subjects' satisfaction with the overview and no-overview interfaces. The overview interface scored significantly higher than the no-overview interface on the dimensions Terrible-Wonderful, $F(1, 28) = 13.81$, $p < .001$; Frustrating-Satisfying, $F(1, 28) = 5.73$, $p < .05$; Rigid-Flexible, $F(1, 28) = 6.73$, $p < .05$; and Keeping track of objects was: Difficult-Easy, $F(1, 28) = 14.45$, $p < .001$. Between maps, we found a significant difference for four satisfaction questions, showing that subjects gave the interfaces higher satisfaction scores after using the Washington map.

4.3 Task Completion Time

Figure 7, panel a, shows the task completion time with the two interfaces and on the two maps. The Washington map was faster overall compared to the Montana

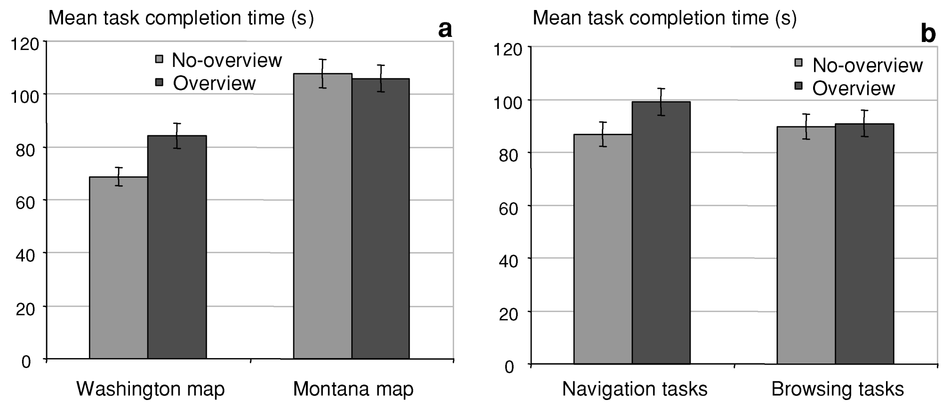

Fig. 7. *Task completion time in seconds* (*for each bar* $N = 160$). This figure shows the mean task completion time in seconds for each solution to a task. Error bars show the standard error of the mean. Panel a shows the task completion times for the Washington and Montana maps. Panel b shows the task completion times for navigation and browsing tasks.

map, $F(1, 28) = 48.94$, $p < .001$. We found a significant interaction between interface and map, $F(1, 28) = 4.50$, $p < .05$. Tasks solved with the no-overview interface on the Washington map were solved 22% faster ($M = 68.76, SD = 43.38$) than tasks solved with the overview ($M = 84.23, SD = 59.42$). Tasks solved on the Montana map were solved with comparable mean completion times (no-overview: $M = 107.81, SD = 68.05$; overview: $M = 105.85, SD = 59.42$).

Going into more detailed analysis using the same analysis of variance as above, we found no significant interaction between task types and interfaces, $F(1, 28) = 1.74$, $p > .1$. However, as can be seen in Figure 7, panel b, the no-overview interface was significantly faster for navigation tasks ($M = 86.9$, $SD = 60.4$), compared to the overview+detail interface ($M = 99.1$, $SD = 64.4$), $F(1, 28) = 5.27$, $p < .05$.

All navigation tasks solved on the Washington map with the no-overview interface had faster task completion times compared to the overview interface. Contradicting our task level hypothesis (hypothesis 3, Section 3.1), we found that one of the navigation tasks that required subjects to compare map objects was solved significantly faster with the no-overview interface (estimated marginal mean = 73.5, $SE = 11.12$) compared to the overview interface (estimated marginal mean = 113.9, $SE = 11.12$), $F(1, 28) = 7.46$, $p < .05$. On the Washington map, four of five browsing tasks were completed faster with the no-overview interface. One of these, a task that required finding the first airport east of some county, was solved significantly faster using the no-overview interface (estimated marginal mean = 81.81, $SE = 11.3$) compared to the overview interface (estimated marginal mean = 122.2, $SE = 11.2$), $F(1, 28) = 5.19$, $p < .05$. This also contradicted our hypothesis.

For the Montana map, no significant differences between interfaces for individual tasks were found. This contradicted our hypotheses that comparison tasks should be performed faster using the overview interface and that browsing tasks should be solved faster using the overview interface.

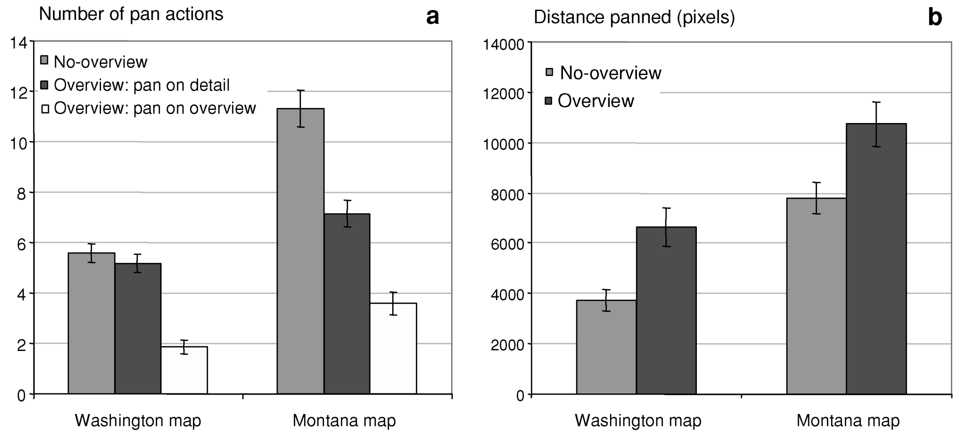

Fig. 8. *Panning in the two interfaces* (*for each bar* $N = 160$). Panel a shows the mean number of pan actions per task in the detail window without overview (left bar) and in the detail window with overview (middle bar), and the panning done by dragging or centering the field-of-view (right bar). Panel b shows the mean distance panned in screen pixels without the overview (left bar) and with the overview (right bar). In both panels, error bars show the standard error of the mean.

Large differences between subjects existed. The slowest subject used on average 169 s/task, or 3.4 times as much as the fastest subject. For individual tasks, differences between subjects were as large as 1:23.

4.4 Navigation Patterns

In the following, we investigate the differences between navigation in the two interfaces and try to provide detailed data about user actions that might explain the differences in task completion time, recall tasks, and satisfaction measures discussed on the preceding pages.

4.4.1 *Number of Pan and Zoom Actions.* Dragging the field-of-view box is the preferred way of panning on the overview. Subjects used this method of panning for half of the tasks solved with the overview. Figure 8, panel a, shows the mean number of panning actions made by panning in the detail view or by centering the field-of-view. We found an interaction effect between map type and interface type, meaning that more pan actions on the detail view happened on the Montana map with the no-overview compared to the overview interface, $F(1, 28) = 12.89$, $p < .05$. However, with the overview, subjects dragged or centered the field-of-view more frequently on the Montana map. Hence, as can be seen in Figure 8, panel b, the overall distance panned, that is, the sum of the distance panned both on the overview and on the detail view, was 51% higher with the overview ($M = 8690$ pixels, $SD = 10{,}554$), compared to no-overview interface ($M = 5{,}751$ pixels, $SD = 6{,}943$), $F(1, 28) = 10.90$, $p < .01$.

In 28% of the tasks solved with the overview, the field-of-view box was resized; in less than 4% of the tasks was the field-of-view box redrawn. Figure 9, panel a, summarizes the zoom actions made by resizing the field-of-view. We found a significant interaction between interface and map type, $F(1, 28) = 19.65$, $p < .001$, meaning that a comparable number of zoom actions were done in the

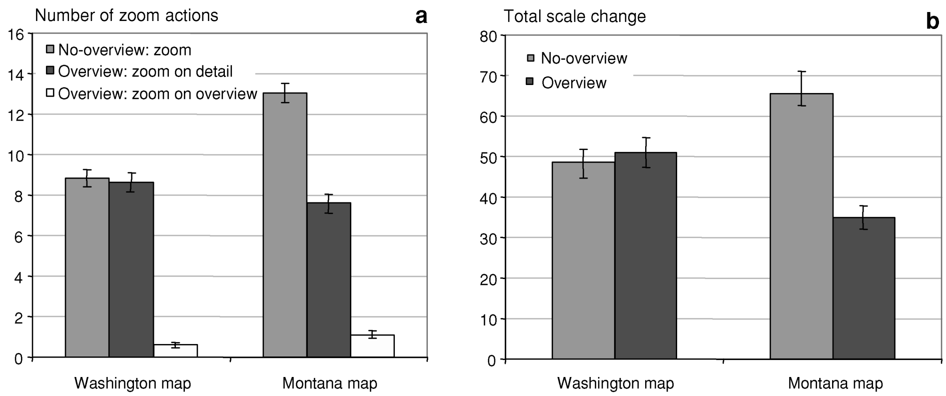

Fig. 9. *Zooming in the two interfaces (for each bar $N = 160$).* Panel a shows the mean number of zoom actions per task in the detail window without overview (left bar) and in the detail window with overview (middle bar), and the zooming done by resizing or redrawing the field-of-view (right bar). Panel b shows the mean scale change without the overview (left bar) and with the overview (right bar). In both panels, error bars show the standard error of the mean.

two interfaces on the Washington map, but that on the Montana map twice as much zooming happened with the no-overview interface as with the overview interface. Subjects seldom zoomed by changing the field-of-view box compared to how often they zoomed on the detail view. Looking at the sum of changes in scale (Figure 9, panel b), we found a significant interaction between interface and map type, $F(1, 28) = 12.70$, $p < .001$. On the Montana map, the no-overview interface ($M = 57$ scales, $SD = 58.9$) has a 33% higher number of scale changes than the overview interface ($M = 43$ scales, $SD = 43.2$), $F(1, 28) = 7.10$, $p < .05$.

4.4.2 *Use of the Overview Window.* In 55% of the 320 tasks solved with the overview, subjects actively interacted with the overview window, that is, they moved or resized the field-of-view box. Tasks in which the overview window was used were frequently solved by first interacting with the detail view then switching to navigating using the overview and then possibly back to the detail view. To better understand the benefit of the overview window, we compared the tasks that were solved by actively using the overview window with the tasks solved without using the overview. Tasks solved with active use of the overview were solved 20% slower (estimated marginal mean = 103.93, $SE = 3.98$) than tasks where the overview window was not actively used (estimated marginal mean = 86.32, $SE = 4.57$), $t(319) = 2.91$, $p < .01$. Another way of understanding the use of the overview window is to look at the transitions between the overview and the detail window. We found that the number of transitions is strongly correlated with task completion time, Spearman's $r = .404$, $p < .001$. The more transitions between the overview and the detail window, the longer the task completion time.

Two subjects did not use the overview at all, while three subjects used the overview at least once for all 10 tasks solved with the overview+detail interface.

4.4.3 *Observations from the Experiment.* We use our notes from observations during the experiment to make three points. First, many subjects

experienced occasional problems with the combined zoom and pan button. Even though subjects practiced this combination button during the training tasks, 18 subjects zoomed at least one time when they verbally indicated that they were trying to pan. The delay before zooming began was sometimes too short. This appeared to happen when subjects began initiating a pan action without having made up their minds about which direction to pan.

Second, subjects' habit formation highlighted some limitations in the interfaces. At least eight subjects tried to use a way of navigating from the overview window in the detail window or vice versa. Some subjects tried to click on the detail window, probably with the intention of jumping to the place where they clicked. This way of navigating seemed to be taken from the overview window, where clicking on a point centers the field-of-view box on that point. Similarly, some subjects tried to zoom in and out while they had the mouse over the overview window. This way of interacting seemed to be mimicked after the interaction with the detail view.

Third, we observed that at least six subjects repeatedly experienced what has been called *desert fog* [Jul and Furnas 1998], that is, they zoomed or panned into an area of the map that contained no map objects. When we observed desert fog, two of these subjects were using the overview interface, four the no-overview interface.

5. DISCUSSION

5.1 Usability and Navigation Patterns

Subjects preferred the overview interface. Subjects also scored this interface significantly higher on the seven satisfaction questions, and commented that the overview helped them keep track of the current position and that the overview window was useful for navigation. This result confirmed our second hypotheses (see Section 3.1) and is coherent with previous empirical work on overviews [North and Shneiderman 2000; Hornbæk and Frøkjær 2001] and recommendations in the design literature [Plaisant et al. 1995, Shneiderman 1998]. For task completion times, we found that subjects who actively used the overview window were slower than subjects who only used the detail window. Our results are surprising considering previous studies, for example, Beard and Walker [1990] and North and Shneiderman [2000], which found that having an overview led to faster task completion times. However, in the studies by Beard and Walker [1990] and North and Shneiderman [2000], navigation in the detail-only interface was done with scrollbars. Our study shows that a direct manipulation zoomable user interface can in some cases (e.g., with the Washington map) reduce—or even eliminate—the need for a separate overview.

We did not find any support for our third hypothesis about an advantage for the overview interface for certain tasks. On the contrary, when considering the difference between browsing and navigation tasks, our results were similar to those of Hornbæk and Frøkjær [2001]. In both studies, it was demonstrated that a no-overview interface can be significantly faster for navigation tasks than an interface with an overview (here in the case of the Washington map).

In the context of our experiment, we consider four explanations of the difference in task completion time between the overview and the no-overview interfaces. First, the overview might be visually distracting, continuously catching subjects' attention and thus affecting task completion time. While we cannot definitively reject this explanation from the data collected, we note that subjects who did not actively use the overview window achieved task completion times comparable to tasks solved with the zoomable user interface (see Section 4.4.2). The straightforward explanation that, since the interface with an overview presented more information, it took more time to use, is also weakened by this observation. A second explanation of the task completion times suggests that switching between the detail and the overview window required mental effort and time moving the mouse. Our data modestly supports this explanation, since the number of transitions between overview and detail window were positively correlated with task completion time. A third explanation is that navigation on the overview window was coarse and that resizing the field-of-view box could be difficult at low zoom factors. Subjects commented that the overview was hard to resize. In support of those comments, we note that the overview window used in the experiment occupied 256×192 pixels. When a zoom factor of 20 was reached, the field-of-view box was only 13×10 pixels, which was probably hard for most users to resize and move using the mouse. Finally, it is conceivable that users never became competent in effectively using the added complexity of the overview. However, it should be noted that our experiment lasted longer than other experiments, for example, North and Shneiderman [2000], that did find an advantage for overviews.

We also investigated how subjects navigated with and without an overview. Interestingly, subjects only directly used the overview in half of the tasks where an overview was available. This rather low figure might indicate that adding zooming to an interface diminishes the value of the overview for navigation purposes compared to nonzoomable interfaces. Subjects panned 51% longer using the overview interface compared to the no-overview interface. One possible explanation for this large difference might be that the overview window did not support fine-grained navigation (as suggested above) and that subjects therefore had to do additional navigation on the detail view. Our data also shows that subjects made more scale changes in the no-overview interface when searching the single-level map. On the single-scale map, there was less information to help navigation. The difference observed might be one indication that the overview helped both navigation and keeping an overview: a function that subjects in the no-overview condition had to substitute for more zooming.

In summary, we found a trade-off between the two interfaces, with the no-overview interface being faster with the Washington map and the overview interface always leading to higher satisfaction. Our results challenge some of the common criticism of zoomable user interfaces without an overview, for example, that users lose their overview when zooming [Card et al. 1999, p. 634]. We found the two interfaces to be comparable with respect to accuracy; on the Washington map, the no-overview interface was faster than the overview interface. We do not know whether the speed difference observed might diminish when users learn to cope with the complexity of the overview interface.

5.2 Possible Influence of Map Organization on Usability

We found surprisingly large differences in usability and navigation patterns between the Washington and the Montana maps; the presence or absence of an overview also interacted with these differences. Because content (Washington vs. Montana) and number of level organization (single vs. multilevel) were confounded (see Section 3.4), this experiment does not allow us to make claims about the origin of those differences. However we believe that the maps (see Figure 3) were similar in most respects, except for the difference in the number of levels. The Washington and Montana maps were similar with respect to the number of map objects (1,591 vs. 1,540) and the area the state occupied (50% vs. 57% of the initial screen). The information density, measured as the mean distance in pixels from any map object to the nearest map object, was also similar (7.1 vs. 7.8). The similarity of the tasks conducted on the two maps was more difficult to measure precisely as browsing tasks require users to make many decisions and often sent them on sidetracks, but we did our best to choose similar tasks. Therefore we believe that the difference between the two maps consisted mostly of a difference in overall organization (single vs. multilevel). We recognize that a separate experiment is needed to verify this hypothesis.

When using the Washington map, subjects were faster and more accurate, and scored the interface higher on subjective satisfaction measures, irrespective of which interface they used. If the difference could be attributed to the use of multiple levels in the Washington map, the result would be consistent with the literature on landmarks [Vinson 1999], since the top-level landmarks—for example, the labels at the lowest scale on the multilevel map—were visible at all scales. Besides being faster with the Washington map, the no-overview interface also improved recall for map locations, partially confirming our first hypothesis. An explanation might be that the richer navigation cues on the Washington map helped the subjects to concentrate navigation and attention on the detail window, thereby relying less on the overview window. A multilevel map might also be more effective because it provides an implicit overview of the space that users memorize as they navigate the detail view. Only two users were seen experiencing desert fog with the Washington map, versus four with the Montana map. Feeling lost and having to reorient oneself, possibly by using the overview window, might be less common with a map like the Washington map.

5.3 Recommendations for Designers and Further Research

An interpretation of our study with the aim of providing advice for designers of information systems offers three main points:

(1) We recommend that designers closely consider the trade-off in subjective satisfaction and task completion time between providing an overview or not. We expect, in most cases, that an overview should be provided, but this depends on the critical usability parameters in the particular context designed for. A walk-up-and-use kiosk should perhaps aim for high satisfaction, while a navigation system for use in time-sensitive situations could

dispense with the overview if the information space contains rich cues for navigation and if the interface provides a flexible way of zooming.

(2) We believe that interfaces with an overview should eliminate navigation commands that are specific only to the overview window or to the detail window, that is, they should aim at unifying navigation [Raskin 2000]. All zoom and pan actions should therefore be similar across windows.

(3) To obtain the benefit of easy navigation provided by overviews (see Section 2.1), designers should use overviews at least one-sixteenth the size of the detail window (in area). For overviews coupled to a detail view less than the size of one screen or for screens on small devices, the overview might need to be larger to support navigation. For systems where much navigation is expected on the overview, for example, in support of monitoring tasks, a larger overview should be provided. For systems with zoom factors over 20 as used in our system, more usability problems will occur when using the overview, and consequently a larger overview will be necessary.

We propose five areas of further research:

(1) We believe that maps organized with multiple levels are likely to be preferable to single-level maps in terms of accuracy, task completion time, and satisfaction. More research needs to be conducted to confirm this hypothesis.

(2) The method for interacting used in the experiment occasionally caused subjects to zoom instead of pan. Research is needed to find a method for interacting with zoomable user interfaces using a two-dimensional input device that is intuitive and supports habit formation. We have used other interaction techniques ourselves, but picked the present interface because we believed it was easier to use for novices. Ideally, zooming and panning should be allowed to take place in parallel.

(3) Empirical research should explore integrating navigation cues within the detail view. Our observations and subjects' comments suggest that a detail-only interface could include cues about the current zoom factor (e.g., Furnas and Zhang [2000]), cues about the current position in the information space, and aids for avoiding desert fog (e.g., Jul and Furnas [1998]). If such cues are integrated into the detail view, the mental and motor effort associated with shifting to the overview might be reduced, as would the screen real estate lost due to the presence of an overview.

(4) Research should aim at improving the usability of the overview window. Usability might be improved by increasing the size of the overview window or by the use of distorted overview windows, which might give users better control over local navigation without losing the possibility of coarse global navigation. Optional overviews, or space multiplexed overviews, might also provide the navigation benefit without constantly taking up screen real estate. In our study the use of the overview for keeping track of ones position in the information space (as opposed to using the overview for navigation) was only addressed in so far as it influenced usability. The problems users encountered when shifting visual and mental attention to

the overview without interacting with it should be further explored, for example, by using eye tracking.

(5) Future research could investigate in more detail the effect on performance of expertise with the information space and the interface. It seems especially important to know how the satisfaction versus time tradeoff develops as users' expertise grows.

6. CONCLUSIONS

We compared the navigation patterns and usability of a zoomable user interface with and without an overview. Thirty-two subjects spent an average of 1.5 h solving tasks on two maps. Our results suggest a tradeoff between the two interfaces in subjective satisfaction and task completion time. Subjects scored the interface with an overview higher on seven subjective satisfaction questions, and 80% preferred this interface. In contrast, subjects were faster without the overview when using one of the two maps. Our results highlight the influence of the map design on usability. Subjects preferred using the Washington map independently of the interface used; they were also significantly faster at completing tasks on this map. We conjecture that this map was more effective because it was organized in multiple levels that might provide an implicit overview through the use of semantic zooming. We also found large individual differences in subjects' ability to navigate the map, in task completion times, and in accuracy. Based on our work, we recommend that the usability of overviews be improved, as should navigation aids for zoomable user interfaces. A better understanding of visual and mental attention in information visualization interfaces would help better explain the usability trade-off found. Common expectations about difficulties with zoomable user interfaces and the utility of an overview were not confirmed in this study. On the contrary, we found that zoomable interfaces without an overview offer certain benefits compared to interfaces with an overview.

APPENDIX: TASKS USED IN THE EXPERIMENT

Washington map, navigation tasks

1. Which city is closest to the city Colton in Whitman County?
2. Which state park is located north of the city Ione in Pend Oreille County?
3. Which of the following two cities is located most to the north: Shelton in Mason County or Warden in Grant County?
4. Which of the following cities covers the largest area: Sequim in Clallam County, Sumas in Whatcom County, or Deer Park in Spokane County?
5. Which are the two largest parks passed on the railroad going from Westport in Grays Harbor County to Vancouver in Clark County?

Washington map, browsing tasks

1. Which two national parks in Washington are biggest?
2. Find and name two counties in Washington that contain two or more military facilities.

3. Find and name the first airport east of the county Skamania.
4. Which two cities in the counties on the northern border of Washington cover the largest area?
5. Which of the counties on the southern border of Washington contains the most cities?

Montana map, navigation tasks

1. Which city is closest to Baker City in Fallon County (in the eastern part of Montana)?
2. Which city is located west of the city Eureka in Lincoln County (in the northwest part of Montana)?
3. Which of the following two cities is located most to the north: Darby in Ravalli County (western part of Montana) or Columbus in Stillwater County (southern part of Montana)?
4. Which of the following cities in the eastern part of Montana covers the largest area: Wolfpoint in Roosevelt County, Glendive in Dawson County, or Ekalaka in Carter County?
5. Which are the two largest cities on the railroad from the city Wibaux in Wibaux County (eastern part of Montana) to the city Red Lodge in Carbon County (southern part of Montana)?

Montana map, browsing tasks

1. Which two lakes in Montana are biggest?
2. Find and name two counties in Montana that contain at least three airports or airfields.
3. Find and name the first state park east of Furgus County (central Montana).
4. Which two cities in the counties on the northern border of Montana cover the largest area?
5. Which of the counties on the southern border of Montana contains the most cities?

ACKNOWLEDGMENTS

We thank all members of the Human-Computer Interaction Laboratory at the University of Maryland for valuable help and encouragement. Erik Frøkjær's and Ben Shneiderman's thoughtful comments improved the paper substantially. For thorough advice on statistics, we thank Per Settergren Sørensen. Constructive and concrete advice from the anonymous reviewers is also acknowledged. The work was done during a 6-month visit of the first author to the Human-Computer Interaction Laboratory at the University of Maryland.

REFERENCES

AHLBERG, C. AND SHNEIDERMAN, B. 1994. Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In *Proceedings of ACM Conference on Human Factors in Computing Systems* (CHI '94, Boston MA, Apr. 24–28). C. Plaisant, Ed. ACM Press, New York, N.Y., 313–317.

BALDONADO, M. Q. W., WOODRUFF, A., AND KUCHINSKY, A. 2000. Guidelines for using multiple views in information visualization. In *Proocedings of the 5th International Working Conference on Advanced Visual Interfaces* (AVI'2000, Palermo, Italy, May 24–26). L. Tarrantino, Ed. ACM Press, New York, N.Y., 110–119.

BEARD, D. B. AND WALKER, J. Q. 1990. Navigational techniques to improve the display of large two-dimensional spaces. *Behav. Inform. Techn. 9*, 6, 451–466.

BEDERSON, B. B. 2001. PhotoMesa: A zoomable image browser using quantum treemaps and bubblemaps. In *UIST 2001, ACM Symposium on User Interface Software and Technology, CHI Lett. 3*, 2, 71–80.

BEDERSON, B. B. AND BOLTMAN, A. 1999. Does animation help users build mental maps of spatial information. In *Proceedings of IEEE Symposium on Information Visualization* (INFOVIZ'99, San Francisco, Calif., Oct. 24–29). G. Wills and D. Keim, Eds. IEEE Press, New York, N.Y., 28–35.

BEDERSON, B. B. AND HOLLAN, J. D. 1994. Pad++: A zooming graphical interface system. In *Proceedings of the 7th ACM Symposium on User Interface Software and Technology* (UIST'94, Marina del Rey, Calif., Nov. 2–4). P. Szekely, Ed. ACM Press, New York, N.Y., 17–26.

BEDERSON, B. B., HOLLAN, J. D., PERLIN, K., MEYER, J., BACON, D., AND FURNAS, G. W. 1996. Pad++: A zoomable graphical sketchpad for exploring alternate interface physics. *J. Vis. Lang. Comput. 7*, 1, 3–31.

BEDERSON, B. B., MEYER, J., AND GOOD, L. 2000. Jazz: An extensible zoomable user interface graphics toolkit in Java. In *UIST'00, ACM Symposium on User Interface Software and Technology, CHI Lett. 2*, 2, 171–180.

CARD, S. K., MACKINLAY, J. D., AND SHNEIDERMAN, B. 1999. *Readings in Information Visualization.* Morgan Kaufmann, San Francisco, Calif.

CARD, S. K., ROBERTSON, G. G., AND MACKINLAY, J. D. 1991. The information visualizer, an information workspace. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI'91, New Orleans, La, Apr. 27–May 2). S. P. Robertson, G. M. Olson, and J. S. Olson, Eds. ACM Press, New York, N.Y., 181–188.

CARR, D., PLAISANT, C., AND HASEGAWA, H. 1998. Designing a real-time telepathology workstation to mitigate communication delays. *Interac. Comput. 11*, 1, 33–52.

CHEN, C. AND CZERWINSKI, M. P. 2000. Special Issue on Empirical Evaluation of Information Visualizations. *Internat. J. Hum.-Comput. Studies 53*, 5.

CHIN, J. P., DIEHL, V. A., AND NORMAN, K. L. 1988. Development of an instrument for measuring user satisfaction of the human-computer interface. In *Proceeding of the ACM Conference on Human Factors in Computing Systems* (CHI '88, Washington, D.C., May 15–19). E. Soloway, D. Frye, and S. B. Sheppard, Eds. ACM Press, New York, N.Y., 213–218.

COMBS, T. AND BEDERSON, B. B. 1999. Does zooming improve image browsing? In *Proceedings of the ACM Conference on Digital Libraries* (DL '99, Berkeley, Calif., Aug. 11–14). N. Rowe and E. A. Fox, A., Eds. ACM Press, New York, N.Y., 130–137.

DRUIN, A., STEWART, J., PROFT, D., BEDERSON, B., AND HOLLAN, J. D. 1997. KidPad: A design Collaboration between children, technologists, and educators. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI '97, Atlanta, Ga, Mar. 22–27). S. Pemperton, Ed. ACM Press, New York, N.Y., 463–470.

EICK, S. G., STEFFEN, J. L., AND SUMNER, E. E. 1992. Seesoft—a tool for visualizing line oriented software statistics. *IEEE Trans. Softw. Eng. 18*, 11, 957–968.

FRANK, A. U. AND TIMPF, S. 1994. Multiple representations for cartographic objects in a multi-scale tree—an intelligent graphical zoom. *Comput. Graph. 18*, 6, 823–829.

FURNAS, G. W. AND BEDERSON, B. B. 1995. Space-scale diagrams: Understanding multiscale interfaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI '95, Denver, Colo., May 7–11). I. R. Katz, R. Mach, L. Marks, M. B. Rosson, and J. Nielsen, Eds. ACM Press, New York, N.Y., 234–241.

FURNAS, G. W. AND ZHANG, X. 1998. MuSE: A multiscale editor. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology* (UIST '98, San Fransisco, Calif., Nov. 1–4). E. Mynatt and R. Jacob, Eds. ACM Press, New York, N.Y., 107–116.

FURNAS, G. W. AND ZHANG, X. 2000. Illusions of infinity: Feedback for infinite worlds. In *UIST 2000, ACM Symposium on User Interface Software and Technology, CHI Lett. 2*, 2, 237–238.

GHOSH, P. AND SHNEIDERMAN, B. 1999. Zoom-only vs. overview-detail pair: A study in browsing techniques as applied to patient histories. University of Maryland Technical Report CS-TR-4028. ftp://ftp.cs.umd.edu/pub/hcil/Reports-Abstracts-Bibliography/99-12html/99-12.html.

GUO, HUO, ZHANG, WEIWEI, AND WU, JING. 2000. The Effect of zooming speed in a zoomable user interface. *Report from Student HCI Online Research Experiments (SHORE)*, http://otal.umd.edu/SHORE2000/zoom/.

HIGHTOWER, R. R., RING, L. T., HELFMAN, J. I., BEDERSON, B. B., AND HOLLAN, J. D. 1998. Graphical multiscale Web histories: A study of PadPrints. In *Proceedings of the Ninth ACM Conference on Hypertext* (Hypertext '98, Pittsburgh, Pa., June 20–24). ACM Press, New York, N.Y., 58–65.

HORNBÆK, K. AND FRØKJÆR, E. 1999. Do thematic maps improve information retrieval? In *IFIP TC.13 International Conference on Human-Computer Interaction* (INTERACT '99, Edingburgh, Scotland, Aug. 30–Sep. 3). M. A. Sasse and C. Johnson, Eds. IOS Press, Amsterdam, The Netherlands, 179–186.

HORNBÆK, K. AND FRØKJÆR, E. 2001. Reading electronic documents: The usability of linear, fisheye, and overview+detail interfaces. In *CHI 2001, ACM Conference on Human Factors in Computing Systems, CHI Lett. 3*, 1, 293–300.

IGARISHI, T. AND HINCKLEY, K. 2000. Speed-dependent automatic zooming for browsing large documents. In *UIST 2000, ACM Symposium on User Interface Software and Technology, CHI Lett. 2*, 2, 139–148.

JUL, S. AND FURNAS, G. W. 1998. Critical zones in desert fog: Aids to multiscale navigation. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology* (UIST '98, San Fransisco, Calif., Nov. 1–4). E. Mynatt and R. Jacob, Eds. ACM Press, New York, N.Y., 97–106.

NORTH, C. AND SHNEIDERMAN, B. 2000. Snap-together visualization: evaluating coordination usage and construction. *Internat. J. Hum.-Comput. Studies, 53*, 5, 715–739.

NORTH, C., SHNEIDERMAN, B., AND PLAISANT, C. 1995. User controlled overviews of an image library: A case study of the visible human. In *Proceedings of the 1st ACM International Conference on Digital Libraries* (DL '96, Bethesda, Md., Mar. 20–23). E. A. Fox and G. Marchionini, Eds. ACM Press, New York, N.Y., 74–82.

PÁEZ, L. B., DA SILVA-FH., J. B., AND MARCHIONINI, G. 1996. Disorientation in electronic environments: A study of hypertext and continuous zooming interfaces. In *Proceedings of the 59th Annual Meeting of the American Society for Information Science* (ASIS '96, Baltimore, Md, Oct. 19–24). S. Harding, Ed., 58–66.

PERLIN, K. AND FOX, D. 1993. Pad: An alternative approach to the computer interface. In *Proceedings of the 20th Annual ACM Conference on Computer Graphics* (SIGGRAPH '93, Anaheim, Calif., Aug. 2–6). J. T. Kajiya, Ed. ACM Press, New York, N.Y., 57–64.

PLAISANT, C., CARR, D., AND SHNEIDERMAN, B. 1995. Image browsers: Taxonomy, guidelines, and informal specifications. *IEEE Softw. 12*, 2, 21–32.

PLAISANT, C. M. B., ROSE, A., AND SHNEIDERMAN, B. 1996. Life lines: Visualizing personal histories. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI '96, Vancouver, B. C., Canada, Apr. 13–18). B., Nardi and G. C., van der Veer, Eds. ACM Press, New York, N.Y., 221–227.

POOK, S., LECOLINET, E., VAYSSEIX, G., AND BARILLOT, E. 2000. Context and interaction in zoomable user interfaces. In *Proocedings of the 5th International Working Conference on Advanced Visual Interfaces* (AVI 2000, Palermo, Italy, May 23–26). L. Tarrantino, Ed. ACM Press, New York, N.Y., 227–231.

RASKIN, J. 2000. *The Humane Interface: New Directions for Designing Interactive Systems.* Addison-Wesley, Reading, Mass.

SCHAFFER, D., ZUO, Z., GREENBERG, S., BARTRAM, L., DILL, J., DUBS, S., AND ROSEMAN, M. 1996. Navigating hierarchically clustered networks through fisheye and full-zoom methods. *ACM Trans. Comput.-Hum. Interact. 3*, 2, 162–188.

SHNEIDERMAN, B. 1998. *Designing the User Interface.* Addison-Wesley, Reading, Mass.

VINSON, N. G. 1999. Design guidelines for landmarks to support navigation in virtual environments. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI '99, Pittsburg, Pa, May 15–20). M. G., Williams, M. W., Altom, K., Ehrlich, and W. Newman, Eds. ACM Press, New York, N.Y., 278–285.

WARE, C. 2000. *Information Visualization: Perception for Design.* Morgan Kaufmann, San Fransisco, Calif.

WOODRUFF, A., LANDAY, J., AND STONEBREAKER, M. 1998a. Constant information density in zoomable interfaces. In *Proocedings of the 4th International Working Conference on Advanced Visual Interfaces* (AVI '98, L'Aquila, Italy, May 24–27). T. Catarci, M. F. Costabile, G. Santucci, and L. Tarantino, Eds. 110–119.

WOODRUFF, A., LANDAY, J., AND STONEBREAKER, M. 1998b. Goal-directed zoom. In *Summary of the ACM Conference on Human Factors in Computing Systems* (CHI '98, Los Angeles, Calif., Apr. 18–23). C.-M. Karat, A. Lund, J. Coutaz, and J. Karat, Eds. ACM Press, New York, N.Y., 305–306.

Received XXX; revised XXX; accepted XXX